

VISIBILITY GRAPHS AND DEFORMATIONS OF ASSOCIAHEDRA

SATYAN L. DEVADOSS, RAHUL SHAH, XUANCHENG SHAO, AND EZRA WINSTON

ABSTRACT. The associahedron is a convex polytope whose face poset is based on nonintersecting diagonals of a convex polygon. In this paper, given an arbitrary simple polygon P , we construct a polytopal complex analogous to the associahedron based on convex diagonalizations of P . We describe topological properties of this complex and provide realizations based on secondary polytopes. Moreover, using the visibility graph of P , a deformation space of polygons is created which encapsulates substructures of the associahedron.

1. ASSOCIAHEDRA FROM POLYGONS

1.1. Let P be a simple planar polygon with labeled vertices. Unless mentioned otherwise, assume the vertices of P to be in general position, with no three collinear vertices. A *diagonal* of P is a line segment connecting two vertices of P which is contained in the interior of P . A *diagonalization* of P is a partition of P into smaller polygons using noncrossing diagonals of P . Let a *convex diagonalization* of P be one which divides P into smaller convex polygons.

Definition 1. Let $\pi(P)$ be the poset of all convex diagonalizations of P where for $a \prec a'$ if a is obtained from a' by adding new diagonals.¹

Figure 1 shows diagonalizations of a polygon P . Parts (b) through (d) show some elements of $\pi(P)$ where part (c) is greater than (d) in the poset ordering.




FIGURE 1. A polygon with (a) a diagonalization and (b) – (d) convex diagonalizations.

If P is a convex polygon, all of its diagonalizations will obviously be convex. The question was asked in combinatorics as to whether there exists a convex polytope whose face poset is isomorphic to $\pi(P)$ for a convex polygon P . It was independently proven in the affirmative by Lee [8] and Haiman (unpublished):

2000 *Mathematics Subject Classification.* Primary 52B11.

Key words and phrases. visibility graph, associahedron, secondary polytope.

¹Mention of *diagonals* will henceforth mean noncrossing ones.

Theorem 2. *When P is a convex polygon with n sides, the associahedron \mathcal{K}_n is a convex polytope of $\dim n - 3$ whose face poset is isomorphic to $\pi(P)$.*

Almost twenty years before this result was discovered, the associahedron had originally been defined by Stasheff for use in homotopy theory in connection with associativity properties of H -spaces [12]. Associahedra have continued to appear in a vast number of mathematical fields, currently leading to numerous generalizations (see [4] [5] for some viewpoints). Figure 2 shows the three-dimensional polyhedron \mathcal{K}_6 on the left. The 1-skeleton of \mathcal{K}_6 is shown on the right, with its vertices labeled with the appropriate diagonalizations of the convex hexagon.

Remark. Classically, the associahedron is based on all bracketings of $n - 1$ letters and denoted as K_{n-1} . Since there is a bijection between bracketings on $n - 1$ letters and diagonalizations of n -gons, we use the script notation \mathcal{K}_n with an index shift to denote the associahedron K_{n-1} for ease of notation in our polygonal context.




FIGURE 2. The 3D associahedron \mathcal{K}_6 .

1.2. We now extend Theorem 2 for arbitrary simple polygons P . A *polytopal complex* S is a finite collection of convex polytopes (containing all the faces of its polytopes) such that the intersection of any two of its polytopes is a (possibly empty) common face of each of them. The dimension of the complex S denoted as $\dim(S)$ is the largest dimension of a polytope in S .

Theorem 3. *For a polygon P with n vertices, there exists a polytopal complex \mathcal{K}_P whose face poset is isomorphic to $\pi(P)$. Moreover, \mathcal{K}_P is a subcomplex of the associahedron \mathcal{K}_n .*

Proof. Let p_1, \dots, p_n be the vertices of P labeled cyclically. For a convex n -gon Q , let q_1, \dots, q_n be its vertices again with clockwise labeling. The natural mapping from P to Q (taking p_i to q_i) induces an injective map $\phi : \pi(P) \longrightarrow \pi(Q)$. Assign to $t \in \pi(P)$ the face of \mathcal{K}_n that corresponds to $\phi(t) \in \pi(Q)$. It is trivial to see that $\phi(t_1) \prec \phi(t_2)$ in $\pi(Q)$ if $t_1 \prec t_2$ in $\pi(P)$.

Moreover, for any $\phi(t)$ in $\pi(Q)$ and any diagonal (q_i, q_j) which does not cross the diagonals of $\phi(t)$, we see that (p_i, p_j) does not cross any diagonal of t . So if a face f of \mathcal{K}_n is contained in a face corresponding to $\phi(t)$, then there exists a diagonalization $t' \in \pi(P)$ where $\phi(t')$ corresponds to f and $t' \prec t$. Since the addition of any noncrossing diagonals to a convex diagonalization is still a convex diagonalization, the intersection of any two faces² is also a face in \mathcal{K}_P . So \mathcal{K}_P satisfies the requirements of a polytopal complex and (due to the map ϕ) is a subcomplex of \mathcal{K}_n . \square

Corollary 4. *Let P be an n -gon and let $d(P)$ be the minimum number of diagonals required to diagonalize P into convex polygons. The polytopal complex \mathcal{K}_P has dimension $n - 3 - d(P)$.*

Proof. The dimension of a polytopal complex is defined as the maximum dimension of any face. In the associahedron \mathcal{K}_n , a face of dimension k corresponds to a convex diagonalization with $n - 3 - k$ diagonals. The result follows since ϕ is an injection. \square

Example. Figure 3 shows two polytopal complexes \mathcal{K}_P for the respective polygons P given. The left side is the 3-dimensional associahedron \mathcal{K}_6 based on diagonalizations of a convex hexagon, whereas the right side is the 2D polytopal complex of a deformed hexagon, made from two edges glued to opposite vertices of a square. Note how this complex appears as a subcomplex of \mathcal{K}_6 .




FIGURE 3. Two polytopal complexes \mathcal{K}_P based on the respective polygons P given.

The diagram given in Figure 4 shows the labeling of the complex given on the right side of Figure 3. Note the number of diagonals in each diagonalization is constant across the dimensions of the faces.

1.3. An alternate construction of \mathcal{K}_P comes from *removing* certain faces of \mathcal{K}_n : Each *facet* of \mathcal{K}_n corresponds to a diagonal of Q . Now consider the set $d_{Q \setminus P}$ of diagonals of Q which are not diagonals of P . If a facet f of \mathcal{K}_n corresponds to a diagonal of $d_{Q \setminus P}$, remove f along with the interior of any face g where $f \cap g$ is nonempty. Notice that this deletes every face that does not correspond to a convex diagonalization of P while preserving all faces that do. It is easy to see that we are left with a polytopal complex \mathcal{K}_P . Since for every two intersecting diagonals there

²Such an intersection could possibly be empty if diagonals are crossing.




FIGURE 4. The labeling of the face poset structure of \mathcal{K}_P given on the right side of Figure 3.

is a third diagonal intersecting neither, any two facets are separated by at most one facet. We therefore have the following:

Lemma 5. *The subcomplex of \mathcal{K}_n which is removed to form \mathcal{K}_P is connected.*

Although the complement of \mathcal{K}_P in \mathcal{K}_n is connected, it is not immediate that \mathcal{K}_P itself is connected. Consider a diagonalization of a polygon P . An *edge flip* of P (called *flip* for short) removes a diagonal of P and replaces it with another noncrossing diagonal of P . The *flip graph* of P is a graph whose nodes are the set of triangulations of P , where two nodes x and y are connected by an arc if one diagonal of x can be flipped to obtain y . It is a classical result of computational geometry that the flip graph of any polygon is connected; see [2] for an overview. Since the flip graph of P is simply the 1-skeleton of \mathcal{K}_P , we obtain:

Theorem 6. *\mathcal{K}_P is connected for any P .*

The right side of Figure 2 along with Figure 4 display flip graphs of the appropriate polygons.

2. TOPOLOGICAL PROPERTIES

2.1. We begin this section by considering arbitrary (not just convex) diagonalizations of P and the resulting geometry of \mathcal{K}_P . Let $\Delta = \{d_1, \dots, d_k\}$ be a set of noncrossing diagonals of P , and let $\mathcal{K}_P(\Delta)$ be the collection of faces in \mathcal{K}_P corresponding to all diagonalizations of P containing Δ .

Lemma 7. *$\mathcal{K}_P(\Delta)$ is a polytopal complex.*

Proof. If t is a diagonalization of P containing Δ , then any t' in $\pi(P)$ must also contain Δ if $t' \prec t$. Thus there must be a face in $\mathcal{K}_P(\Delta)$ that corresponds to t' . Furthermore, consider faces f_1 and f_2 in $\mathcal{K}_P(\Delta)$ corresponding to diagonalizations t_1 and t_2 of P . Then the intersection of f_1 and f_2 must correspond to a diagonalization including Δ since $f_1 \cap f_2$ is the (possibly empty) face corresponding to all convex diagonalizations that include every diagonal of t_1 and t_2 . \square

The following lemma is an immediate consequence of the construction of $\mathcal{K}_P(\Delta)$. It shows the gluing map between two faces of \mathcal{K}_P .

Lemma 8. *Let Δ_1 and Δ_2 be two collections of noncrossing diagonals of P . Then $\mathcal{K}_P(\Delta_1)$ and $\mathcal{K}_P(\Delta_2)$ are glued together in \mathcal{K}_P along the (possibly empty) polytopal subcomplex $\mathcal{K}_P(\Delta_1 \cup \Delta_2)$.*

Theorem 9. *Suppose the diagonals $\Delta = \{d_1, \dots, d_k\}$ divide P into polygons Q_0, \dots, Q_k . Then $\mathcal{K}_P(\Delta)$ is isomorphic to the cartesian product $\mathcal{K}_{Q_0} \times \dots \times \mathcal{K}_{Q_k}$.*

Proof. We use induction on k . When $k = 1$, any face $f \in \mathcal{K}_P(d)$ corresponds³ to a convex diagonalization of Q_0 paired with a convex diagonalization of Q_1 . Thus, a face of $\mathcal{K}_{Q_0} \times \mathcal{K}_{Q_1}$ exists for each pair of faces (f_0, f_1) , for $f_0 \in \mathcal{K}_{Q_0}$ and $f_1 \in \mathcal{K}_{Q_1}$. For $k > 1$, order the diagonals such that d_k divides P into polygons $Q_* = Q_0 \cup \dots \cup Q_{k-1}$ and Q_k . A face in $\mathcal{K}_P(\Delta)$ corresponds to a convex diagonalization t_1 of Q_* and a convex diagonalization t_2 of Q_k . The pair $(t_1, t_2) \in \mathcal{K}_P(\Delta)$ corresponds to a face in $\mathcal{K}_{Q_*}(\Delta \setminus d_k) \times \mathcal{K}_{Q_k}$. By the induction hypothesis, $\mathcal{K}_{Q_*}(\Delta \setminus d_k)$ is isomorphic to $\mathcal{K}_{Q_0} \times \mathcal{K}_{Q_1} \times \dots \times \mathcal{K}_{Q_{k-1}}$. \square

Corollary 10. *If the diagonals Δ divide P into convex polygons with Q_0, \dots, Q_k where Q_i has n_i edges, then $\mathcal{K}_P(\Delta)$ is the product of associahedra $\mathcal{K}_{n_0} \times \dots \times \mathcal{K}_{n_k}$.*

Thus every face of \mathcal{K}_P is a product of associahedra. For a polytopal complex \mathcal{K}_P , the *maximal* elements of its face poset $\pi(P)$ are analogous to facets of convex polytopes. These elements of $\pi(P)$ are characterized by the following:

Definition 11. A face f of \mathcal{K}_P corresponding to a diagonalization $t \in \pi(P)$ is a *maximal face* if there does not exist $t' \in \pi(P)$ such that $t \prec t'$.

Thus a maximal face of \mathcal{K}_P has a convex diagonalization of P using the *minimal* number of diagonals. Figure 5 shows a polygon P along with six minimal convex diagonalizations of P . As this shows, such diagonalizations may not necessarily have the same number of diagonals.




FIGURE 5. Six minimal convex diagonalizations of a polygon.

Corollary 10 shows that each maximal face is a product of associahedra. Moreover, by Lemma 8, the maximal faces of \mathcal{K}_P can be glued together to construct \mathcal{K}_P . Figure 6 shows an example of \mathcal{K}_P for the polygon shown. It is a polyhedral subcomplex of the 5-dimensional convex associahedron \mathcal{K}_8 . We see that \mathcal{K}_P is made of six maximal faces, four squares (where each square is a product

³We abuse notation by writing $\mathcal{K}_P(d) = \mathcal{K}_P(\{d\})$.

$\mathcal{K}_4 \times \mathcal{K}_4$ of line segments) and two \mathcal{K}_6 associahedra. Each of these six faces correspond to the minimal convex diagonalizations of P given in Figure 5.




FIGURE 6. An example of \mathcal{K}_P seen as a 3-dimensional subcomplex of the 5-dimensional \mathcal{K}_8 . It is made of six maximal faces, four $\mathcal{K}_4 \times \mathcal{K}_4$ squares and two \mathcal{K}_6 associahedra.

2.2. Having already shown that \mathcal{K}_P is connected, we now prove that this polytopal complex is indeed contractible. We start with some basic geometry: A vertex of a polygon is called *reflex* if the diagonal between its two adjacent vertices cannot exist. Note that every nonconvex polygon has a reflex vertex.

Lemma 12. *For any reflex vertex v of a nonconvex polygon P , every element of $\pi(P)$ has at least one diagonal incident to v .*

Proof. Assume otherwise and consider an element of $\pi(P)$. In this convex diagonalization, since there is no diagonal incident to v , there exists a unique subpolygon containing v . Since v is reflex, this subpolygon cannot be convex, which is a contradiction. \square

Lemma 13. *Let $F = \{f_1, f_2, \dots, f_k\}$ be a set of faces of \mathcal{K}_P such that $\bigcap_F f_i$ is nonempty. If $\bigcap_{F'} f_i$ is contractible for every $F' \subset F$, then $\bigcup_F f_i$ is contractible.*

Proof. We prove this by induction on the number of faces in F . A single face is trivially contractible. Now assume $\bigcup_F f_i$ is contractible. For a face $f_{k+1} \notin F$ of \mathcal{K}_P , let $G = \{f_{k+1}\} \cup F$ so that $\bigcap_G f_i$

is nonempty. Since f_{k+1} intersects the intersection of the faces F and since this intersection is nonempty and contractible, we can deformation retract $\bigcup_G f_i$ onto $\bigcup_F f_i$ and hence maintain contractibility. \square

Theorem 14. *For any polygon P , the polytopal complex \mathcal{K}_P is contractible.*

Proof. We prove this by induction on the number of vertices. For the base case, note that \mathcal{K}_P is a point for any triangle P . Now let P be a polygon with n vertices. If P is convex, then \mathcal{K}_P is the associahedron \mathcal{K}_n and we are done. For P nonconvex, let v be a reflex vertex of P . Since each diagonal d of P incident to v separates P into two smaller polygons Q_1 and Q_2 , by our hypothesis, \mathcal{K}_{Q_1} and \mathcal{K}_{Q_2} are contractible. Theorem 9 shows that $\mathcal{K}_P(d)$ is isomorphic to $\mathcal{K}_{Q_1} \times \mathcal{K}_{Q_2}$, resulting in $\mathcal{K}_P(d)$ to be contractible.

Let $\Delta = \{d_1, \dots, d_k\}$ be the set of all diagonals incident to v . Since Δ is a set of *noncrossing* diagonals, then $\bigcap_{\Delta} \mathcal{K}_P(d)$ is nonempty. Furthermore, for any subset $\Delta' \subset \Delta$, we have $\mathcal{K}_P(\Delta') = \bigcap_{\Delta'} \mathcal{K}_P(d)$. By Theorem 9, this is a product of contractible pieces, and thus itself is contractible. Therefore, by Lemma 13, the union of the complexes $\mathcal{K}_P(d_i)$ is contractible. However, since Lemma 12 shows that this union is indeed \mathcal{K}_P , we are done. \square

Remark. A recent paper by Braun and Ehrenborg [3] considers a similar complex for nonconvex polygons P . Their complex $\theta(P)$ is a *simplicial* complex with vertex set the diagonals in P and facets given by triangulations of P . Indeed, it is easy to see that $\theta(P)$ is the combinatorial dual of \mathcal{K}_P . Figure 7 shows an example of $\theta(P)$ on the left and \mathcal{K}_P on the right, as depicted in Figure 4. For




FIGURE 7. The simplicial complex $\theta(P)$ and its corresponding dual \mathcal{K}_P .

a nonconvex polygon P with n vertices, the central result in [3] shows that $\theta(P)$ is homeomorphic to a ball of dimension $n - 4$. This is, in fact, analogous to Theorem 14 showing \mathcal{K}_P is contractible. The method used in [3] is based on a *pairing lemma* of Linusson and Shareshian [9], motivated by discrete Morse theory, whereas our method considers the geometry based on reflex vertices.

2.3. The previous constructions and arguments can be extended to include planar polygons with holes. These *generalized polygons* are bounded, connected planar regions whose boundary is the disjoint union of simple polygonal loops. Vertices, edges, and diagonals are defined analogously to that for polygons. In what follows, denote R as a generalized polygon with n vertices and $h + 1$ boundary components. Similar to Definition 1, one can define $\pi(R)$ to be the set of all

convex diagonalizations of R partially ordered by inclusion. We state the following and leave the straight-forward proof⁴ to the reader:

Corollary 15. *There exists a polytopal complex \mathcal{K}_R whose face poset is isomorphic with $\pi(R)$. Moreover, the dimension of \mathcal{K}_R is $n + 3h - d(R) - 3$, where $d(R)$ is the minimum number of diagonals required to diagonalize R into convex polygons.*

Moreover, since there is at least one reflex vertex in any generalized polygon (with more than one boundary component), the proof of the following is identical to that of Theorem 14:

Corollary 16. *\mathcal{K}_R is contractible for a generalized polygon R . In particular, \mathcal{K}_R is connected.*

Figure 8 shows an example of the associahedron of a pentagon with a triangular hole, whose maximal faces are 8 cubes and 3 squares. Similar to the polygonal case, the complexes \mathcal{K}_R can




FIGURE 8. The associahedron for a pentagon with a triangular hole, along with its 11 maximal faces.

be obtained by gluing together different polytopal complexes \mathcal{K}_P , for various polygons P . For simplicity, we now describe this in the case when R has only two boundary components.

Let Δ be the set of diagonals of R connecting one vertex of the outside boundary region of R with one vertex of the inside hole of R . Observe that for any diagonal d in Δ , there exists a nonconvex polygon P such that \mathcal{K}_P is isomorphic to $\mathcal{K}_R(d)$. The reason is that there is a bijection between all the convex diagonalizations of R containing d and all the convex diagonalizations of P obtained by splitting the diagonal d into two edges in the polygonal boundary, as given in Figure 9.

⁴Recall the classical result that the number of diagonals in a triangulation of R is $n + 3h - 3$.

Since any convex diagonalization of R contains at least one diagonal in Δ , then \mathcal{K}_R is constructed from gluing elements of $\{\mathcal{K}_R(d) \mid d \in \Delta\}$. Now if the number of boundary components of R is more than two, a similar argument as above can be applied to reduce the number of boundary components by one, and induction will yield the result.




FIGURE 9. A diagonal of a polygon with a hole leads to a nonconvex polygon without a hole.

Remark. In order for an n -gon with a k -gon hole to *not* have a polyhedral complex which is a subcomplex of \mathcal{K}_{n+k+2} , every diagonal from the hole to the exterior polygon must be optional (i.e. not required in every triangulation). Otherwise, the construction given in Figure 9 can be used to find such a map. In particular, the complex given in Figure 8 is not a subcomplex of \mathcal{K}_{10} . Given any generalized polygon R , it is an interesting question to find the smallest n such that \mathcal{K}_R is a subcomplex of \mathcal{K}_n .

3. GEOMETRIC REALIZATIONS

3.1. This section focuses on providing a geometric realization of \mathcal{K}_P with integer coordinates for its vertices. There are numerous realizations of the classical associahedron [10] and its generalizations [11] [5]. We first recall the method followed in [6] for the construction of \mathcal{K}_n , altered slightly to conform to this paper. This is introduced in order to contrast it with the secondary polytope construction which follows.

Let P be a polygon and choose an edge e of P to be its base. Any triangulation T of P has a dual structure as a rooted tree, where the root of the tree corresponds to the edge e . For a triangle Δ of T , let $r(\Delta)$ be the collection of triangles of T that have a path to the root (in the dual tree) containing Δ . Figure 10 gives three examples of the values $|r(\Delta)|$ for triangles in different triangulations of polygon; the triangle adjacent to the root is shaded.

Define a value $\theta_T(\Delta)$ for each triangle of T using the recursive condition

$$(3.1) \quad \sum_{x \in r(\Delta)} \theta_T(x) = 3^{|r(\Delta)|-1}.$$

For each vertex p_i of P , assign the value

$$\hat{\theta}_T(p_i) = \max \{ \theta_T(\Delta) \mid p_i \in \Delta \}.$$





FIGURE 10. Assigning values to triangles in different triangulations of a polygon.

Theorem 17. [6, Section 2] Let p_1, \dots, p_n be vertices of a convex polygon P such that p_{n-1} and p_n are the two vertices incident to the root edge e of P . The convex hull⁵ of the points

$$(\hat{\theta}_T(p_1), \dots, \hat{\theta}_T(p_{n-2}))$$

in \mathbb{R}^{n-2} , as T ranges over all triangulations of P , yields the associahedron \mathcal{K}_n .

Notice that since both of the vertices adjacent to the root edge will have the same value $\hat{\theta}$ for all triangulations, we simply ignore them. Moreover, it is easy to see from Eq. (3.1) that this realization of \mathcal{K}_n must lie in the hyperplane $\sum x_i = 3^{n-3}$ of \mathbb{R}^{n-2} . Figure 11 shows the values $\hat{\theta}$ given to each vertex in each triangulation, along with the convex hull of the five points.

FIGURE 11. An example of Theorem 17 producing \mathcal{K}_5 .

⁵Given a finite set of points S in \mathbb{R}^n , we define the *convex hull* of S to be the smallest convex set containing S , whereas the *hull* of S is the boundary of the convex hull.

Remark. For nonconvex polygons, since \mathcal{K}_P is a subcomplex of \mathcal{K}_n by Theorem 3, the method above extends to a realization of \mathcal{K}_P with integer coordinaters.

3.2. The notion of a secondary polytope was developed by Gelfand, Kapranov, and Zelevinsky [7]. We review it and show its applicability to \mathcal{K}_P . Let P be a polygon with vertices p_1, \dots, p_n . For a triangulation T of P , let

$$\phi(p_i) = \sum_{p_i \in \Delta \in T} \text{area}(\Delta)$$

be the sum of the areas of all triangles Δ which contain the vertex p_i . Let the *area vector* of T be

$$\Phi(T) = (\phi(p_1), \dots, \phi(p_n)).$$

Definition 18. The *secondary polytope* of $\Sigma(P)$ of a polygon P is the convex hull of the area vectors of all triangulations of P .

The following is a special case of the work on secondary polytopes:

Theorem 19. [7, Chapter 7] *If P is a convex n -gon, the secondary polytope $\Sigma(P)$ is a realization of the associahedron \mathcal{K}_n .*

It is easy to see that all the area vectors of P lie in an $(n-3)$ -dimensional plane of \mathbb{R}^n . However, it is not immediate that each area vector is on the hull of $\Sigma(P)$.

We show that the secondary polytope of a nonconvex polygon has all its area vectors on its hull, as is the case for a convex polygon. However, since the secondary polytope of a nonconvex polygon is *not* a subcomplex of the secondary polytope for convex polygon, this result is not trivial.⁶

Theorem 20. *For any polygon P , all area vectors $\Phi(T)$ lie on the hull of $\Sigma(P)$.*

Proof. Fix a triangulation T of P . We first show that there is a height function ω on P which raises the vertices of T to a locally convex surface in \mathbb{R}^3 , that is, a surface which is convex on every line segment in P . Choose an edge e of P to be its base so that the dual tree of T is rooted at e . Starting from the root and moving outward, assign increasing numbers m_i to each consecutive triangle Δ_i in the tree. Define a height function


$$\omega(p_i) = \min \{m_k \mid p_i \in \Delta_k\}$$

for each vertex p_i of P . Observe that for every pair of adjacent triangles Δ_1 and Δ_2 (in the dual tree), we can choose the value m_i to be large enough such that the planes containing $\omega(\Delta_1)$ and $\omega(\Delta_2)$ are distinct and meet in a convex angle; see Figure 12 below.

Now, in order to show that $\Phi(T)$ lies on the hull of $\Sigma(P)$, we construct a linear function $\rho(v)$ on $\Sigma(P)$ such that $\rho(\Phi(T))$ is a unique minimum of this function on $\Sigma(P)$. For any v in $\Sigma(P)$, define $\rho(v) = \langle \omega(T), v \rangle$ to be the inner product of the vectors $v \in \mathbb{R}^n$ and

$$\omega(T) = (\omega(p_1), \dots, \omega(p_n)).$$

⁶The original proof for the convex case uses a *regular* triangulation observation; we modify this using a lifting map of the dual tree of the polygon.

FIGURE 12. A height function ω lifting the triangles of T .

For a triangle Δ of T with vertices p_i, p_j, p_k , consider the volume in \mathbb{R}^3 enclosed between Δ and the lifted triangle $\omega(\Delta)$. This volume can be written as

$$\frac{\omega(p_i) + \omega(p_j) + \omega(p_k)}{3} \text{area}(\Delta).$$

The volume between the surface on which the $\omega(p_i)$'s lie and the plane is given by

$$\begin{aligned} \sum_{\Delta \in T} \frac{\omega(p_i) + \omega(p_j) + \omega(p_k)}{3} \text{area}(\Delta) &= \sum_{i=1}^n \left[\frac{\omega(p_i)}{3} \sum_{p_i \in \Delta \in T} \text{area}(\Delta) \right] \\ &= \sum_{i=1}^n \frac{\omega(p_i)}{3} \phi(p_i) \\ &= \frac{1}{3} \langle \omega(T), \Phi(T) \rangle. \end{aligned}$$

Since ω lifts T to a locally convex surface S , we know that w will lift any $T' \neq T$ to a surface S' above S . Thus $\langle \omega(T), \Phi(T) \rangle < \langle \omega(T), \Phi(T') \rangle$, implying all vertices of $\Sigma(P)$ lie on the hull. \square

Corollary 21. *If P is nonconvex, then a subset of the faces of $\Sigma(P)$ yield a realization of \mathcal{K}_P .*

Proof. For any face f of \mathcal{K}_P , let T_1, \dots, T_k be the triangulations corresponding to the vertices of f . We use the same argument as the theorem above to show there exists a height function ω such that $\langle \omega, \Phi(T) \rangle$ is constant for any $T \in \{T_1, \dots, T_k\}$ and $\langle \omega, \Phi(T) \rangle < \langle \omega, \Phi(T') \rangle$ for any $T' \notin \{T_1, \dots, T_k\}$. \square

4. VISIBILITY GRAPHS

4.1. We now consider the space of simple planar polygons through deformations. Throughout this section, we only consider simple polygons with vertices labeled $\{1, 2, \dots, n\}$ in this cyclic order.

As before, we assume the vertices of P to be in general position, with no three collinear vertices. Let us begin with the notion of visibility.

Definition 22. The *visibility graph* $\mathcal{V}(P)$ of a labeled polygon P is the labeled graph with the same vertex set as P , with e as an edge of $\mathcal{V}(P)$ if e is an edge or diagonal of P .

A classical open problem in visibility of polygons is as follows:

Open Problem. Given a graph G , find nice necessary and sufficient conditions which show if there exists a simple polygon P such that $\mathcal{V}(P) = G$.

Definition 23. Two polygons P_1 and P_2 are \mathcal{V} -equivalent if $\mathcal{V}(P_1) = \mathcal{V}(P_2)$.

There is a natural relationship between the graph $\mathcal{V}(P)$ and the polytopal complex \mathcal{K}_P : If polygons P_1 and P_2 are \mathcal{V} -equivalent then \mathcal{K}_{P_1} and \mathcal{K}_{P_2} yield the same complex. We wish to classify polygons under a stronger relationship than \mathcal{V} -equivalence. For a polygon P , let (x_i, y_i) be the coordinate of its i -th vertex in \mathbb{R}^2 . We associate a point $\gamma(P)$ in \mathbb{R}^{2n} to P where

$$\gamma(P) = (x_1, y_1, x_2, y_2, \dots, x_n, y_n).$$

Since P is labeled, it is obvious that γ is injective but not surjective.

Definition 24. Two polygons P_1 and P_2 are \mathcal{V} -isotopic if there exists a continuous map $f : [0, 1] \rightarrow \mathbb{R}^{2n}$ such that $f(0) = \gamma(P_1)$, $f(1) = \gamma(P_2)$, and for every $t \in [0, 1]$, $f(t) = \gamma(P)$ for some simple polygon P where $\mathcal{V}(P) = \mathcal{V}(P_1)$.

It follows from the definition that two polygons that are \mathcal{V} -isotopic are \mathcal{V} -equivalent. The converse is not necessarily true: Figure 13 shows three polygons along with their respective visibility graphs. Parts (a) and (c) are \mathcal{V} -equivalent (having identical visibility graphs) but cannot be deformed into one another without changing their underlying visibility graphs. The middle figure (b) shows an intermediate step in obtaining a deformation.




FIGURE 13. Polygons (a) and (c) are \mathcal{V} -equivalent but not \mathcal{V} -isotopic. The visibility graphs for each of the polygons are shown.

For a polygon P with n vertices, let $\mathcal{D}(P)$ be the \mathcal{V} -isotopic equivalence class containing the polygon P and let \mathcal{D} be the set of all such equivalence classes of polygons with n vertices. We give \mathcal{D} a poset structure: For two n -gons P_1 and P_2 , the relation $\mathcal{D}(P_2) \prec \mathcal{D}(P_1)$ is given if the following two conditions hold:

- (1) $\mathcal{V}(P_1)$ is obtained by adding one more edge to $\mathcal{V}(P_2)$.
- (2) There exists a continuous map $f : [0, 1] \longrightarrow \mathbb{R}^{2n}$, such that $f(0) = \gamma(P_1)$, $f(1) = \gamma(P_2)$, and for every $t \in [0, 1/2)$, $f(t) = \gamma(P)$ for some polygon P with $\mathcal{V}(P) = \mathcal{V}(P_1)$, while for every $t \in (1/2, 1]$, $f(t) = \gamma(Q)$ for some polygon Q with $\mathcal{V}(Q) = \mathcal{V}(P_2)$.

If P_1 and P_2 are \mathcal{V} -isotopic, let $\mathcal{D}(P_1) = \mathcal{D}(P_2)$. Taking the transitive closure of \preceq yields the *deformation poset* \mathcal{D} . A natural ranking exists on \mathcal{D} based on the number of edges of the visibility graphs.

Example. Figure 14 shows a subdiagram of the Hasse diagram for \mathcal{D} for 6-gons. For the sake of presentation, we have forgone the labeling on the vertices. A polygonal representative for each equivalence class is drawn, along with its underlying visibility graph. Each element of \mathcal{D} corresponds to a complex \mathcal{K}_P as displayed in Figure 15. Notice that as the polygon deforms and loses visibility edges, its associated polytopal complex collapses into a vertex of \mathcal{K}_6 .

4.2. It is easy to see that the deformation poset \mathcal{D} is connected: Notice that \mathcal{D} has a unique *maximum* element corresponding to the convex polygon. Given any polygon P in the plane, one can move its vertices, deforming P into convex position, making each element of \mathcal{D} connected to the maximum element. Since the vertices of P are in general position, we can insure that the visibility graph of P changes only by one diagonal at a time during this deformation. However, during this process, the visibility graph of the deforming polygon might gain and lose edges, moving up and down the poset \mathcal{D} .


We are interested in the combinatorial structure of the deformation poset beyond connectivity. The maximum element of \mathcal{D} corresponds to the convex n -gon (with $\binom{n}{2}$ edges in its visibility graph) whereas the minimal elements (which are not unique in \mathcal{D}) correspond to polygons with unique triangulations (with $2n - 3$ edges in each of their visibility graphs). This implies that the height of the deformation poset is

$$\binom{n}{2} - 2n + 4.$$


We pose the following problem and close this paper with a discussion of partial results:

Deformation Problem. *Show that every maximal chain of \mathcal{D} has length $\binom{n}{2} - 2n + 4$.*

We can rephrase this open problem loosely in the deformation context: Does there exist a deformation of any simple polygon into a convex polygon such that throughout the deformation the visibility of the polygon monotonically increases? And moreover, does there exist a deformation of any simple polygon into a polygon with a unique triangulation such that throughout the deformation the visibility of the polygon monotonically decreases?

FIGURE 14. A subdiagram of the Hasse diagram of \mathcal{D} for 6-gons.

For the remaining part of the paper, we focus on the monotonically increasing segment of the deformation problem. A natural approach is to discretize this problem into moving vertices of the polygon one by one. So a stronger claim is as follows: For any (nonconvex) polygon, there exists a vertex which can be moved in the plane that preserves the visibility of vertices and introduces a new visibility. In other words, for any polygon, there exists a vertex which can be moved such that we can always move up in rank in the poset. Figure 16 provides an elegant counterexample to this claim [1], obtained by the participants of the first Mexican workshop on Computational Geometry (<http://xochitl.matem.unam.mx/~rfabila/DF08/>). A partial collection of the visibility edges

FIGURE 15. The subcomplexes of \mathcal{K}_6 associated to the elements of Figure 14.

of this polygon is given in red. For this polygon, no vertex can be moved to increase visibility without first losing its current visibility.

We do have the following positive result in the special case of *star* polygons. A polygon P is a star polygon if there exists a point $p \in P$ such that p is visible to all points of P .

Theorem 25. *Let P be a star n -gon. There exists a chain in \mathcal{D} from P to the maximum element.*

Proof. Let x be a point in the *kernel* of P , the set of points which are visible to all points of P . Choose an ε -neighbourhood around x contained in the kernel. For any $a \in P$, let $p(a)$ be a point on the boundary of P which is the intersection of the ray from x passing through a with the




FIGURE 16. No vertex may be moved to increase visibility without first decreasing visibility.

boundary. Let

$$r(a) = \frac{d(a, x)}{d(p(a), x)}$$

and let a' be the point on the ray from x passing through a such that $d(a', x) = \varepsilon \cdot r(a)$. Let ϕ be the map from a to a' . We thus construct a linear map $f : P \times [0, 1] \rightarrow P$ where $f(P, 0) = P$ and $f(P, 1) = \phi(P)$ and where

$$\frac{\partial}{\partial t} f(a, t) = r(a).$$

For any two visible vertices a and b of P , consider the triangle abx . There cannot be any vertices of P contained in the triangle. If for any vertex c of P , the ray from x passing through c intersects the line segment ab at a point z , then $d(c, x) > d(z, x)$ and thus for no $t \in [0, 1)$ can $d(f(c, t), x) \leq d(f(z, t), x)$. So no visibility is lost during the transformation, but notice that $\phi(P)$ is a circle. However, if we apply $f(a, t)$ only to the vertices of P and map any point z on an edge (a, b) of P to z' on the edge between $f(a, t)$ and $f(b, t)$, we find that we get a polygon at every t . Moreover, the edge is always further from c than $f(z, t)$ for every $t \in [0, 1]$, and thus visibility is still maintained. \square

REFERENCES

1. O. Aichholzer, personal communication.
2. P. Bose and F. Hurtado. Flips in planar graphs, *Computational Geometry: Theory and Apps.* **42** (2009) 60-80.
3. B. Braun and R. Ehrenborg, The complex of noncrossing diagonals of a polygon, preprint arXiv:0802.1320.

4. M. Carr and S. Devadoss. Coxeter complexes and graph-associahedra, *Topology and its Appl.* **153** (2006) 2155-2168.
5. F. Chapoton, S. Fomin, A. Zelevinsky. Polytopal realizations of generalized associahedra, *Can. Math. Bull.* **45** (2002) 537-566.
6. S. L. Devadoss, A realization of graph associahedra, *Disc. Math.* **309** (2009) 271-276.
7. I.M. Gelfand, M.M. Kapranov and A.V. Zelevinsky, *Discriminants, resultants, and multidimensional determinants*, Birkhäuser, Boston (1994).
8. C. Lee, *The associahedron and triangulations of the n -gon*, *European J. Combin.* **10** (1989) 551-560.
9. S. Linusson and J. Shareshian, Complexes of t -colorable graphs, *SIAM J. Disc. Math.* **16** (2003) 371-389.
10. J.-L. Loday. Realization of the Stasheff polytope, *Archiv der Mathematik* **83** (2004) 267-278.
11. A. Postnikov. Permutohedra, associahedra, and beyond, *Int. Math. Res. Not.*, to appear.
12. J. Stasheff, Homotopy associativity of H -spaces I, *Trans. Amer. Math. Soc.* **108** (1963) 275-292.

S. DEVADOSS: WILLIAMS COLLEGE, WILLIAMSTOWN, MA 01267
E-mail address: `satyan.devadoss@williams.edu`

R. SHAH: WILLIAMS COLLEGE, WILLIAMSTOWN, MA 01267
E-mail address: `rahul.a.shah@williams.edu`

X. SHAO: MIT, CAMBRIDGE, MA 02139
E-mail address: `zero@mit.edu`

E. WINSTON: BARD COLLEGE, ANNANDALE-ON-HUDSON, NY 12504
E-mail address: `ew429@bard.edu`